

The Evolution of Cyberinfrastructure for Science: from IP to OIL

William E. Johnston

<http://www-itg.lbl.gov/~wej/>

***Computational Research Division,
DOE Lawrence Berkeley National Lab***

and

***NASA Advanced Supercomputing Division,
NASA Ames Research Center***

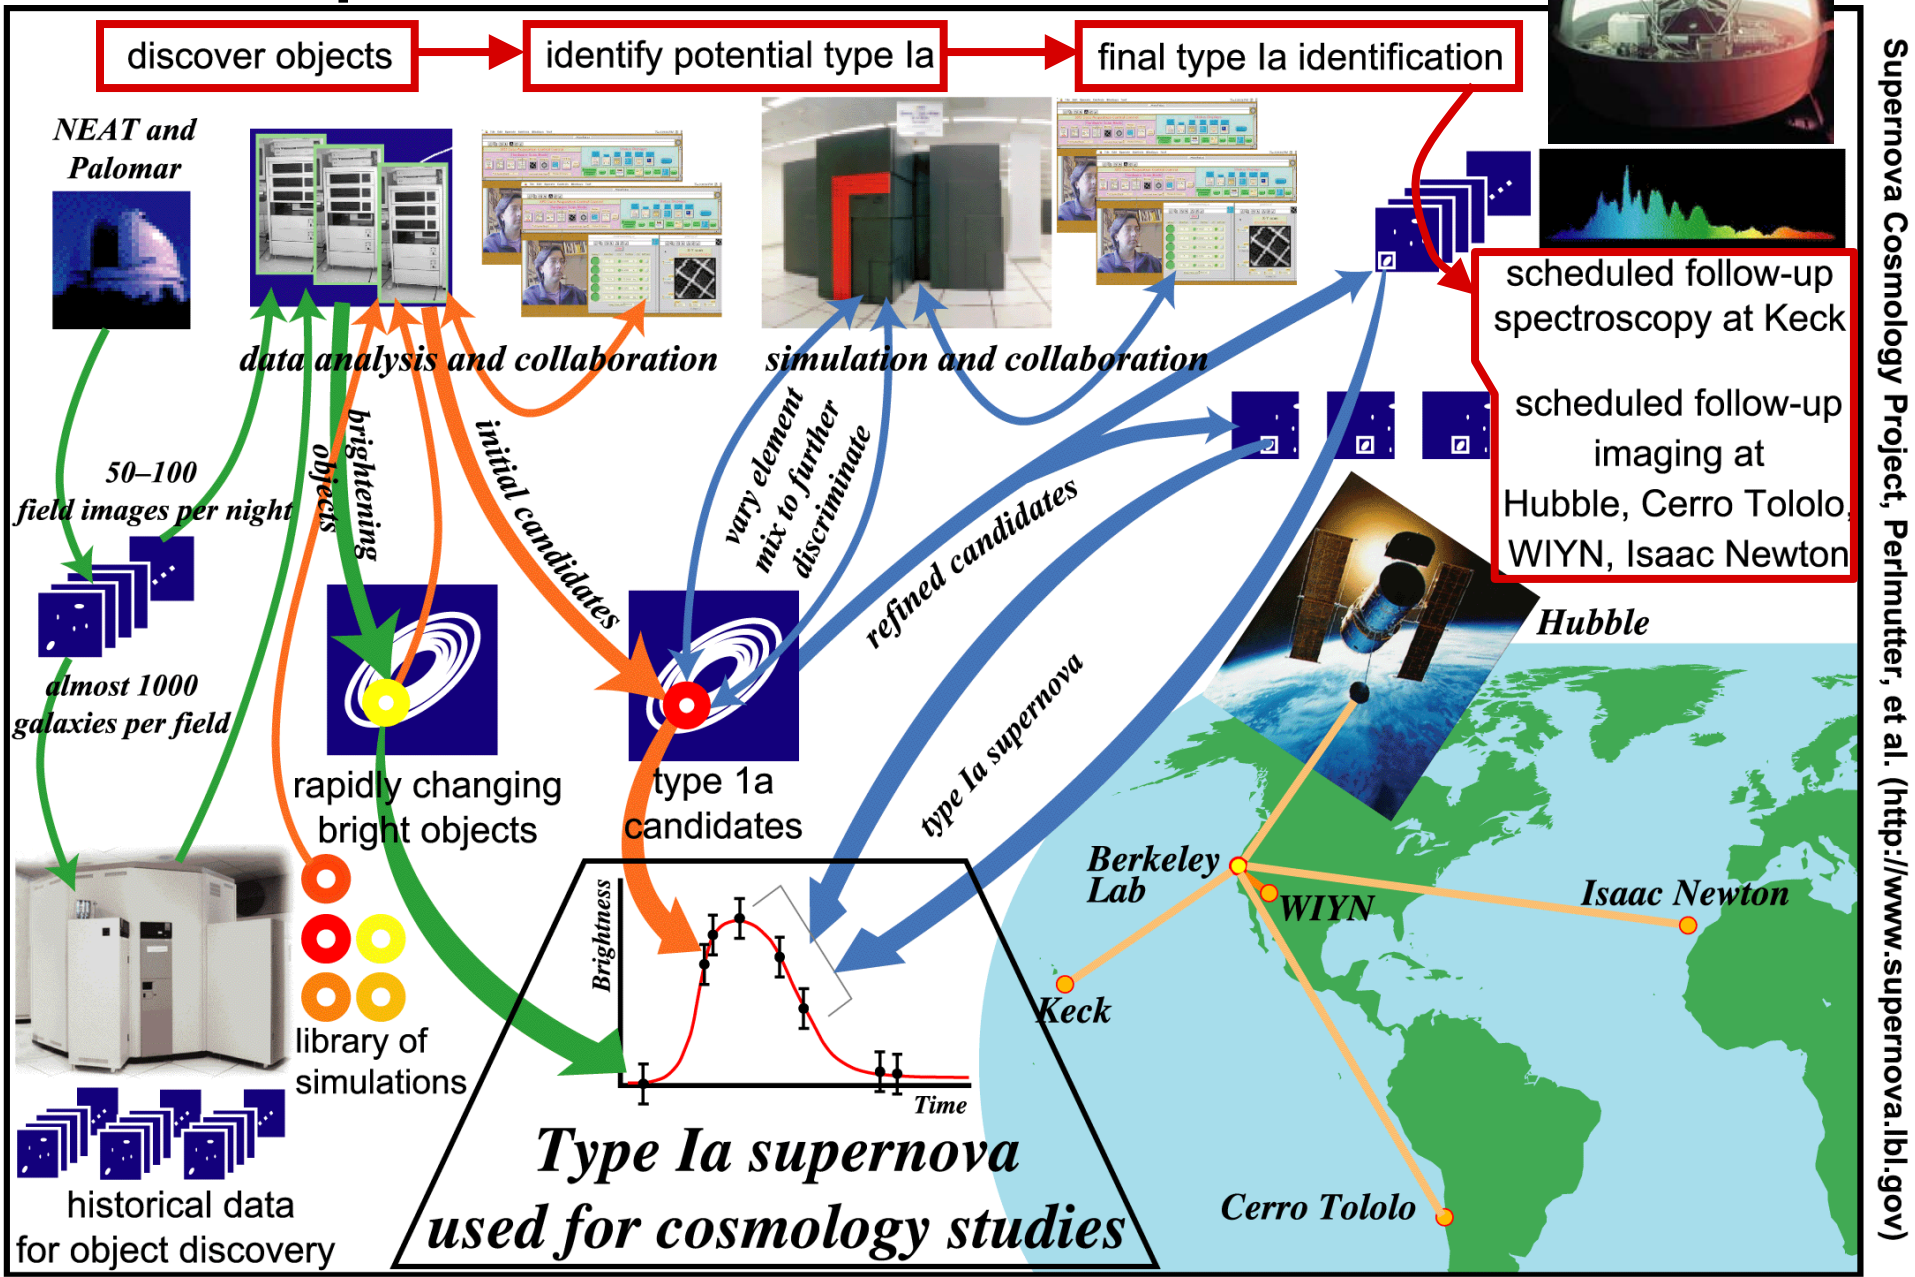
The Process of Large-Scale Science is Changing

- Distributed data, computing, people, instruments
- Instruments integrated with large-scale computing and data systems
- In some fields – e.g. astronomy and astrophysics – science is being done by mining data from dozens of instruments, instead of the direct use of a single instrument like a telescope

Change Drives the Evolution of Cyberinfrastructure

- Large-scale science and engineering problems *require collaborative use of many compute and data resources*, including supercomputers and large-scale data storage systems, *all of which must be integrated with applications and data* that are
 - developed by different teams of researchers
 - or that are obtained from different instrumentsand all of which are at different geographic locations.

Supernova Cosmology is a Complex, Distributed Collaboration



Cyberinfrastructure

- Such complex scenarios require sophisticated infrastructure – high-speed networks, very high-speed computers, and highly capable middleware – to support the application frameworks that are needed to successfully describe and manage the many operations needed to carry out the science.

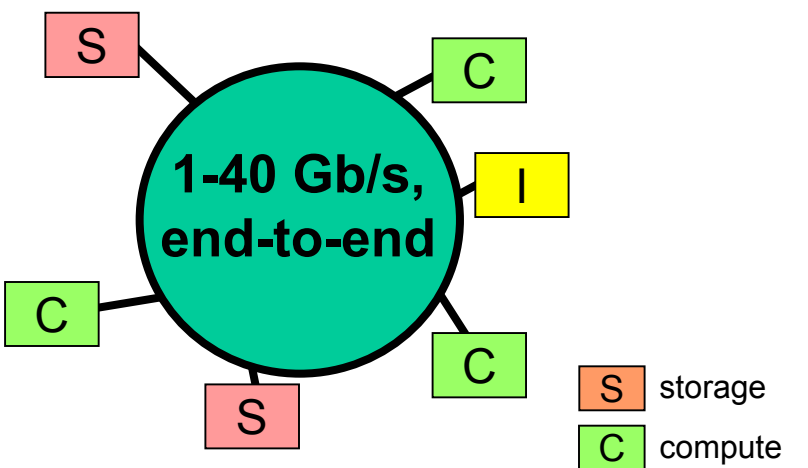
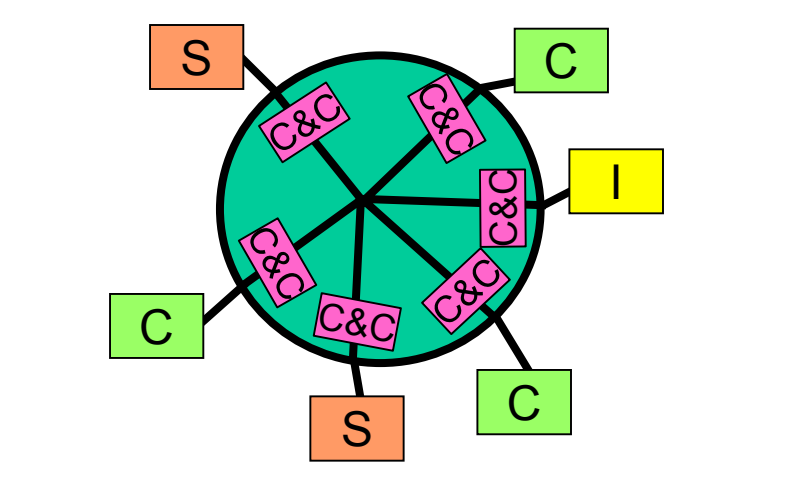
Cyberinfrastructure

- In order to inform the development and deployment of technology, *a set of high-impact science applications in the areas of high energy physics, climate, chemical sciences, magnetic fusion energy, and molecular biology have been analyzed** to characterize their visions for the future process of science, and the networking and middleware capabilities needed to support those visions

*DOE Office of Science, High Performance Network Planning Workshop. August 13-15, 2002: Reston, Virginia, USA.

<http://doecollaboratory.pnl.gov/meetings/hnpnw>

Evolving Requirements for Network Related Infrastructure

1-3 yrs	 <p>1-40 Gb/s, end-to-end</p> <p>S storage C compute I instrument C&C cache & compute</p>	2-4 yrs
In the near term applications need high bandwidth	The next requirement is for bandwidth and QoS.	
3-5 yrs		4-7 yrs
The next requirement is for bandwidth and QoS and network resident cache and compute elements.		The next requirement is for bandwidth and QoS and network resident cache and compute elements, and robust bandwidth.

Evolving Requirements for Middleware

Capabilities to support scientists / engineers / domain problem solvers

<ul style="list-style-type: none"> –Collaboration tools (work group management, document sharing and distributed authoring, sharing application session, human communication) –Programmable portals - facilities to express, manipulate, preserve the representation of scientific problem solving steps (AVS, MatLab, Excel, SciRun) 	<ul style="list-style-type: none"> –Data discovery (“super SQL” for globally distributed data repositories), management, mining, cataloging, publish –Human interfaces (PDAs, Web clients, high-end graphics workstations) –Tools to build/manage dynamic virtual organizations –Knowledge management
---	---

Capabilities to support building the portals / frameworks / problem solving environments

<ul style="list-style-type: none"> –Resource discovery, brokering, job management –Workflow management –Grid management – fault detection and correction 	<ul style="list-style-type: none"> –Grid monitoring and information distribution – event publish and subscribe –Security and authorization,
---	---

Capabilities to support instantiating science scenarios as computational models

<ul style="list-style-type: none"> –Utilities for visualization, data management (global naming, location transparency (replication mgmt, caching), metadata management, data duration, discovery mechanisms) 	<ul style="list-style-type: none"> –Support for programming on the Grid (Grid MPI, Globus I/O, Grid debuggers, programming environments, e.g., to support the model coupling frameworks, and to), Grid program execution environment –User services (documentation, training, evangelism)
--	---

Capabilities to support building Grid systems

Current Cyberinfrastructure

technology	impact
Internet Protocol (IP) – transport independent of the type of the underlying network. Routing (how do I get where I want to go?) Domain Name System (basic directory service: lbl.gov ->128.3.7.82)	The Internet gave us a basic global cyberinfrastructure for transport and interprocess communication.
Secure Socket Layer (SSL)/ Transport Layer Security (TLS)	The ability to communicate securely between known and authenticated endpoints.
Hyper Text Markup Language (HTML) Standardized, low-level document formatting	The Web gave us an information bazaar
Grids Computational and Data Grid services provide us with a global infrastructure for dynamically managing and securely accessing compute, data, instrument, and human resources	The Grid is providing the infrastructure for large-scale collaborative science.

Evolution of Cyberinfrastructure

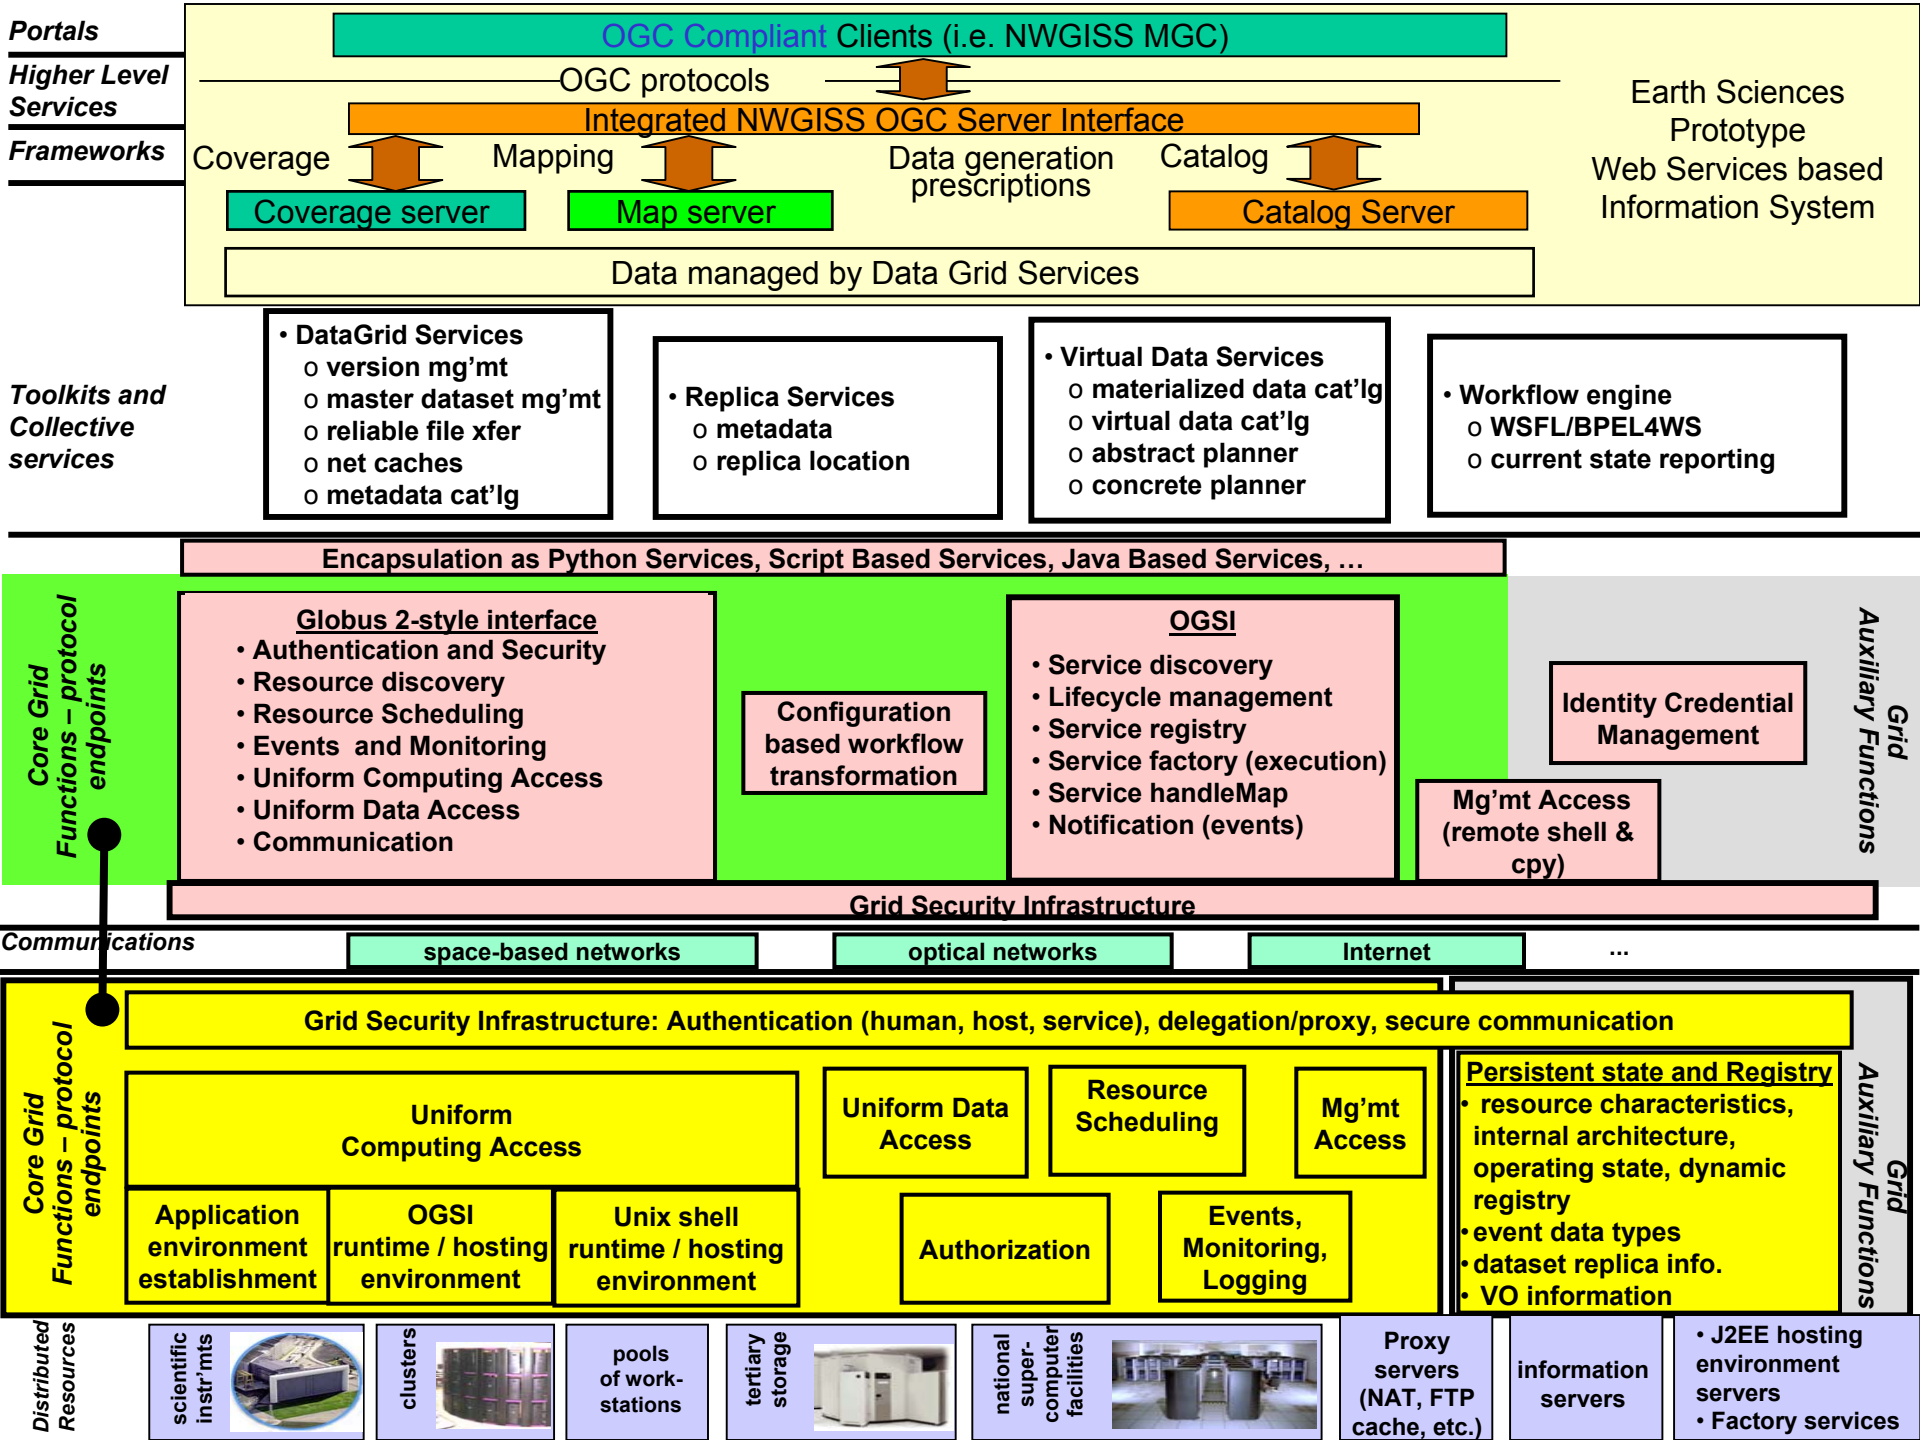
technology	impact
<p>eXtended Markup Language (XML)</p> <p>Tagged fields (metadata) and structured “documents” (XML schema)</p>	<p>Web Services – an application of XML – will provide us with a global infrastructure for managing and accessing modular programs (services) and well defined data</p>
<p>Open Grid Services Architecture</p>	<p>Combines Web Services with Computational and Data Grids services to integrate information, analysis tools, and computing and data systems</p>

Computing and Data Grids

- Grids are several types of middleware that provide
 - computing and data system discovery
 - secure and uniform access to computing and data systems
 - uniform security services
 - tools and services for
 - management of complex workflow involving many compute and data intensive steps occurring at different geographic locations
 - autonomous fault management and recovery for both applications and infrastructure
 - managing large, complex data archives – e.g. geometry (structure) and performance of airframes and turbomachines – that are maintained by discipline experts at different sites, and must be accessed and updated by collaborating scientists
 - distributing and managing massive datasets that must be accessible accessible by world-wide collaborations – e.g. high energy physics
- Grids are also several hundred people from around the world working on best practice and standards at the Global Grid Forum (www.gridforum.org)

Web Services and Grids

- Web services provide for
 - Describing services/programs with sufficient information that they can be discovered and used by many people (reusable components)
 - Assembling groups of discovered services into useful problem solving systems
 - Easy integration with scientific databases that use XML based metadata
- So, Web Services provide for defining, accessing, and managing services, while Grids provide for accessing and managing compute and data systems, and provide support for Virtual Organizations / collaborations



Implications of combining Web Services and Grids

There is considerable potential benefit to combining Grid Services and Web services.

Evolution of Cyberinfrastructure and the Grid

technology	impact
eXtended Markup Language (XML) Tagged fields (metadata) and structured “documents” (XML schema)	Web Services – an application of XML – will provide us with a global infrastructure for managing and accessing modular programs (services) and well defined data
Open Grid Services Architecture	Combines Web Services with Computational and Data Grids services to integrate information, analysis tools, and computing and data systems

Web Services and Grids

- Web services provide for
 - Describing services/programs with sufficient information that they can be discovered and used by many people (reusable components)
 - Assembling groups of discovered services into useful problem solving systems
 - Easy integration with scientific databases that use XML based metadata
- Grids provide for accessing and managing compute and data systems, and provide support for Virtual Organizations / collaborations

The Open Grid Services Architecture (Web + Grid Services)

- From Web services
 - Standard interface definition mechanisms
 - Interface and implementation (multiple protocol bindings)
 - local/remote transparency
 - Language interoperability
 - A homogenous architecture basis
- From Grids
 - Service semantics
 - Lifecycle management and transient state
 - Reliability and security models
 - Discovery
 - Other services: resource management, authorization, etc.
- See <http://www.globus.org/ogsa/>

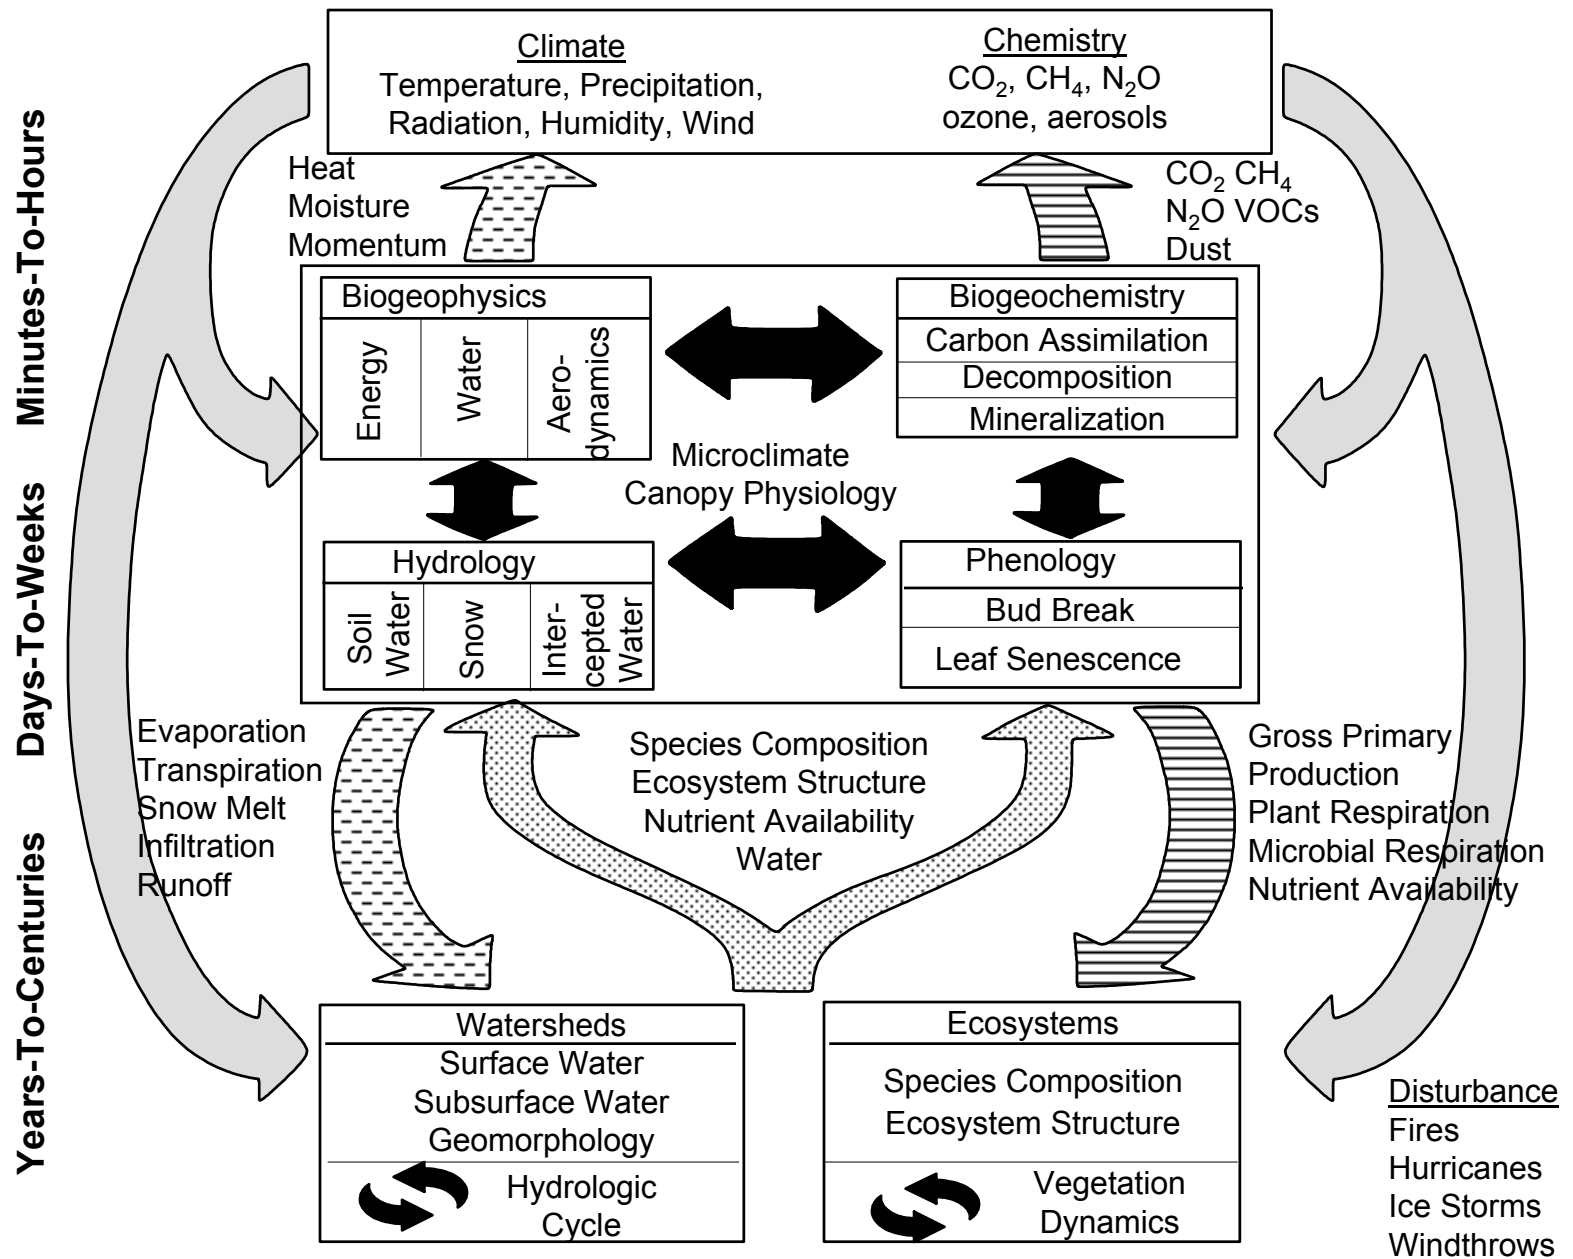
Combining Web Services and Grids

- Combining Grid and Web services provides a dynamic and powerful computing and data environment that is rich in descriptions, services, data, and computing capabilities
- This infrastructure will give us the basic tools to deal with complex, multi-disciplinary, data rich science modeling problems

Combining Web Services and Grids

- Furthermore, the Web Services Description Language, et al, together with Grid services, should be able to provide standardized component descriptions and interface definitions so that compatible sub-models can be “plugged” together
- The complexity of the modeling done in Terrestrial Biogeoscience is a touchstone for this stage of evolution of cyberinfrastructure

Terrestrial Biogeoscience Involves Many Complex Processes and Data



(Courtesy Gordon Bonan, NCAR: *Ecological Climatology: Concepts and Applications*. Cambridge University Press, Cambridge, 2002.)

Where to in the Future?

The Semantic Grid: Beyond Web Services and Grids

- Even when we have well integrated Web+Grid services we still do not provide enough structured information to let us ask “what if” questions, and then have the underlying system assemble the required components in a consistent way to answer such a question.

Beyond Web Services and Grids

- A commercial example “what if” question: What does the itinerary look like if I wish to go SFO to Paris, CDG, and then to Bucharest. In Bucharest I want a 3 or 4 star hotel that is within 3 km of the Palace of the Parliament, and the hotel cost may not exceed the U. S. Dept. of State, Foreign Per Diem Rates.
- To answer such a question – relatively easy, but tedious, for a human – the system must “understand” the relationships between maps and locations, between per diem charts and published hotel rates, and it must be able to apply constraints (< 3 km, 3 or 4 star, $\text{cost} < \$ \text{ per diem rates}$, etc.)

Beyond Web Services and Grids

- A science example (necessarily more complex) that is courtesy of Stewart Loken, LBNL is as follows:

HEP experiments collect specific types of data for the particles that result from high energy collisions of the protons, electrons, ions, etc. that are produced by the accelerators. The types of data are a function of the detector and include things like particle charge, mass, energy, 3D trajectory, etc.

However much of science comes from inferring other aspects of the interactions by analyzing what can be observed. Many quantities are used in obtaining the scientific results of the experiment that are derived from what is observed. In doing this more abstract analysis, the physicist typically asks questions like :

Beyond Web Services and Grids

Events of interest are usually characterized by a combination of jets of particles (coming from quark decays) and single particles like electrons and muons. In addition, we look for missing transverse energy (an apparent failure of momentum conservation) that would signal the presence of neutrinos that we cannot detect.

The topologies of individual events follow some statistical distributions so it is really the averages over many events that are of interest. In doing the analysis, we specify what cone angle would characterize a jet, how far one jet needs to be from another (in 3-dimensions), how far from the single particles, how much missing transverse energy, the angles between the missing energy vector and the other particles

What I would like to see is a set of tools to describe these topologies without typing in lots of code. A graphical interface that lets you draw the average event and trace out how statistical variations would affect that. We do simulation of interesting processes and they guide the selection of events, so we would want to learn from that.

In order to transform these sorts of queries into combinations of existing tools and appropriate data queries, ***some sort of knowledge-based framework is needed.***

Knowledge Grids / Semantic Grids

- The emerging Knowledge Grid* / Semantic Grid** services will provide the mechanisms to organize the information and services so that human queries may be correctly structured for the available application services (the model components and data) to build problem solving systems for specific problems
- Work is being adapted from the Artificial Intelligence community to provide
 - Ontology languages to extend terms (metadata) to represent relationships between them
 - Language constructs to express rule based relationships among, and generalizations of the extended terms
- See www.isi.cs.cnr.it/kgrid/ and www.semanticgrid.org

* I am indebted to Mario Cannataro, Domenico Talia, and Paolo Trunfio (CNR, Italy), and

** Dave DeRoure (U. Southampton) and Carol Gobel (U. Manchester) for introducing me to these ideas.

Future Cyberinfrastructure*

technology	Impact
<p>Resource Description Framework (RDF)**</p> <p>Expresses relationships among “resources” (URI(L)s) in the form of object-attribute-value (property). Values of can be other resources, thus we can describe arbitrary relationships between multiple resources.</p> <p>RFD uses XML for its syntax.</p>	<p>Can ask questions like “What are a particular property’s permitted values, which types of resources can it describe, and what is its relationship to other properties.”</p>
<p>Resource Description Framework Schema (RDFS)**</p> <p>An extensible, object-oriented type system that effectively represents and defines classes.</p>	<p>Object-oriented structure: Class definitions can be derived from multiple superclasses, and property definitions can specify domain and range constraints.</p> <p>Can now represent tree structured information (e.g. Taxonomies)</p>

* See “The Semantic Web and its Languages,” an edited collection of articles in IEEE Intelligent Systems, Nov. Dec. 2000. D. Fensel, editor.

** The Resource Description Framework,” O. Lassila. *ibid.*

Future Cyberinfrastructure

technology	impact
<p>Ontology Inference Layer (OIL)**</p> <p>OIL inherits all of RDFS, and adds expressing class relationships using combinations of intersection (AND), union (OR), and compliment (NOT). Supports concrete data types (integers, strings, etc.)</p>	<p>OIL can state conditions for a class that are both sufficient and necessary. This makes it possible to perform automatic classification: Given a specific object, OIL can automatically decide to which classes the object belongs.</p> <p>This is functionality that should make it possible to ask the sort of constraint and relationship based questions illustrated above.</p>
<p>DAML+OIL+.....***</p>	<p>Knowledge representation and manipulation that have well defined semantics and representation of constraints and rules for reasoning</p>

** “FAQs on OIL: Ontology Inference Layer,” van Harmelen and Horrocks. *ibid.* and “OIL: An Ontology Infrastructure for the Semantic Web.” *Ibid.*


*** “Semantic Web Services,” McIlraith, Son, Zeng. *Ibid.* and “Agents and the Semantic Web,” Hendler. *Ibid.*

Future Cyberinfrastructure


- This Knowledge Grid / Semantic Grid framework should give us the ability to answer “what if” questions by “automatically” structuring data and simulation / analysis components into workflows whose composite actions produce the desired information.
- At the last Grid Forum meeting a Semantic Grid Research Group was established to investigate and report on the path forward for combining Grids and the Semantic Web.

See <http://www.semanticgrid.org/GGF>


This GGF Research Group is co-chaired by David De Roure <dder@ecs.soton.ac.uk>, Carole Goble <cgoble@cs.man.ac.uk>, and Geoffrey Fox <gcf@grids.ucsf.edu>



Supercomputing and Large-Scale Storage



Advanced
Engine
Design




Macromolecular
Crystallography




Advanced
Photon Source



Supernova
Observatory




Advanced
Chemistry



High Energy
Physics




Spallation
Neutron
Source




References

- 1) "The Computing and Data Grid Approach: Infrastructure for Distributed Science Applications," William E. Johnston.
<http://www.itg.lbl.gov/~johnston/Grids/homepage.html#CI2002>
- 2) "Implementing Production Grids for Science and Engineering," William E. Johnston, Lawrence Berkeley National Laboratory, Berkeley, Calif. and NASA Ames Research Center, Moffett Field, Calif. (USA); John M. Brooke, Manchester Computing and Department of Computer Science, University of Manchester (UK); Randy Butler, National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign (USA); David Foster, CERN LHC Computing Grid Project, Geneva (Switzerland), and; Mirco Mazzucato, INFN-Padova (Italy)
<http://www.itg.lbl.gov/~johnston/Grids/homepage.html#Grid2ed>
- 3) NASA's Information Power Grid - www.ipg.nasa.gov
- 4) DOE Science Grid - www.doesciencegrid.org
- 5) European Union DataGrid Project - www.eu-datagrid.org/
- 6) UK eScience Program - www.research-councils.ac.uk/escience/
- 7) NSF TeraGrid - www.teragrid.org/
- 8) GriPhyN (Grid Physics Network) - <http://www.griphyn.org>
- 9) Particle Physics Data Grid, PPDG. <http://www.ppdg.net/>
- 10) European Union DataGrid Project - www.eu-datagrid.org/



Supercomputing and Large-Scale Storage



Advanced
Engine
Design




Macromolecular
Crystallography




Advanced
Photon Source




Spallation
Neutron
Source



Supernova
Observatory



Advanced
Chemistry



High Energy
Physics



Spallation
Neutron
Source

	Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline				Networking	Middleware
Climate (near term)		<ul style="list-style-type: none"> • A few data repositories, many distributed computing sites • NCAR - 20 TBy • NERSC - 40 TBy • ORNL - 40 TBy 		<ul style="list-style-type: none"> • Authenticated data streams for easier site access through firewalls 	<ul style="list-style-type: none"> • Server side data processing (computing and cache embedded in the net) • Information servers for global data catalogues
Climate (5 yr)		<ul style="list-style-type: none"> • Add many simulation elements/components as understanding increases • 100 TBy / 100 yr generated simulation data, 1-5 PBy / yr (just at NCAR) <ul style="list-style-type: none"> – Distribute to major users in large chunks for post-simulation analysis 	<ul style="list-style-type: none"> • Enable the analysis of model data by all of the collaborating community 	<ul style="list-style-type: none"> • Robust access to large quantities of data 	<ul style="list-style-type: none"> • Reliable data/file transfer <ul style="list-style-type: none"> – Across system / network failures
Climate (5+ yr)		<ul style="list-style-type: none"> • 5-10 PBy/yr (at NCAR) • Add many diverse simulation elements/components, including from other disciplines - this must be done with distributed, multidisciplinary simulation • Virtualized data to reduce storage load 	<ul style="list-style-type: none"> • Integrated climate simulation that includes all high-impact factors 	<ul style="list-style-type: none"> • Robust networks supporting distributed simulation - adequate bandwidth and latency for remote analysis and visualization of massive datasets 	<ul style="list-style-type: none"> • Quality of service guarantees for distributed, simulations • Virtual data catalogues and work planners for reconstituting the data on demand

	Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline				Networking	Middleware
High Energy Physics (near term (1-2yr))		<ul style="list-style-type: none"> • Instrument based data sources • Hierarchical data repositories • Hundreds of analysis sites • Petabytes of data 	<ul style="list-style-type: none"> • Productivity aspects of rapid response 	<ul style="list-style-type: none"> • Gigabit/sec • end-to-end QoS 	<ul style="list-style-type: none"> • Secure access to world-wide resources • Data migration in response to usage patterns and network performance <ul style="list-style-type: none"> – naming and location transparency • Deadline scheduling for bulk transfers • Policy based scheduling / brokering for the ensemble of resources needed for a task • Automated planning and prediction to minimized time to complete task
HEP (3-5 yr)		<ul style="list-style-type: none"> • 100s of petabytes of data • Global collaboration • Compute and storage requirements will be satisfied by optimal use of all available resources 	<ul style="list-style-type: none"> • Worldwide collaboration will cooperatively analyze data and contribute to a common knowledge base • Discovery of published (structured) data and its provenance 	<ul style="list-style-type: none"> • 100 Gigabit/sec <ul style="list-style-type: none"> – lambda based point-to-point for single high b/w flows – capacity planning • Network monitoring 	<ul style="list-style-type: none"> • Track world-wide resource usage patterns to maximize utilization • Direct network access to data management systems • Monitoring to enable optimized use of network, compute, and storage resources • Publish / subscribe and global discovery
HEP (5-10 yr)		<ul style="list-style-type: none"> • 1000s of petabytes of data 		<ul style="list-style-type: none"> • 1000 Gigabit/sec 	

Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline			Networking	Middleware
Chem. Sci. (near term)	<ul style="list-style-type: none"> • High data-rate instruments • Greatly increased simulation resolution- data sets ~10 – 30 TB • Geographically separated resources (compute, viz, storage, instmts) & people • Numerical fidelity and repeatability • Cataloguing of data from a large number of instruments 	<ul style="list-style-type: none"> • Distributed collaboration • Remote instrument operation / steering (?) • Remote visualization • Sharing of data and metadata using web-based data services 	<ul style="list-style-type: none"> • Robust connectivity • Reliable data transfer • High data-rate, reliable multicast • QoS • International interoperability for namespace, security 	<ul style="list-style-type: none"> • Collaboration infrastructure • Management of metadata • High data integrity • Global event services • Cross discipline repositories • Network caching • Server side data processing • Virtual production to improve tracability of data • Data Grid broker / planner • Cataloguing as a service
Chem. Sci. (5 yr)	<ul style="list-style-type: none"> • 3D Simulation data sets 30 - 100 TB • Coupling of MPP quantum chemistry and molecular dynamics simulations • Validation using large experimental data sets 	<ul style="list-style-type: none"> • Remote steering of simulation time step • Remote data sub-setting, mining, and visualization • Shared data/metadata w annotation evolves to knowledge base 	<ul style="list-style-type: none"> • 10's Gigabit for collaborative viz and mining of large data sets 	<ul style="list-style-type: none"> • Remote I/O • Collaborative • International interoperability for collab. infrastructure, repositories, search, and notification • Archival publication
Chem. Sci. (5+ yr)	<ul style="list-style-type: none"> • Accumulation of archived simulation feature data and simulation data sets • Multi-physics and soot simulation data sets ~1 PB 	<ul style="list-style-type: none"> • Internationally collab knowledgebase • Remote collaborative simulation steering, mining, viz. 	<ul style="list-style-type: none"> • 100? Gigabit for distributed computation chemistry and molecular dynamics simulations 	<ul style="list-style-type: none"> • Remote collaborative simulation steering, mining, viz

Feature	Characteristics that Motivate H-S Nets	Vision for the Future Process of Science	Anticipated Requirements	
Discipline			Networking	Middleware
Magnetic Fusion (near term)	<ul style="list-style-type: none"> • Each experiment only gets a few days per year - high productivity is critical • 100 MBy every 15 minutes to be delivered in two minutes • Highly collaborative environment 	<ul style="list-style-type: none"> • Real time data analysis for experiment steering (the more that you can analyze between shots the more effective you can make the next shot) 		
Fusion (5 yr)	<ul style="list-style-type: none"> • 1000 MBy generated by experiment every 15 minutes (time between shots) to be delivered in two minutes • 1000 MBy generated by simulation to be delivered in two minutes for comparison with experiment • Simulation data scattered across US • Transparent security • Global directory and naming services needed to anchor all of the distributed metadata • Support for “smooth” collaboration in a high stress environments 	<ul style="list-style-type: none"> • Real time data analysis for experiment steering combined with simulation interaction = big productivity increase • Real time visualization and interaction among collaborators across US • Integrated simulation of the several distinct regions of the reactor will produce a much more realistic model 	<ul style="list-style-type: none"> • 500 Mbit/sec for 20 seconds out of 15 minutes, guaranteed • QoS • 5 to 10 remote sites involved for data analysis and visualization 	<ul style="list-style-type: none"> • Parallel network I/O between simulations, data archives, experiments, and visualization • High quality, 7x24 PKI infrastructure • end-to-end QoS • QoS management • Secure / authenticated transport to ease access through firewalls • Reliable data transfer • transient and transparent data replication for real-time reliability • Collaboration support
Fusion (5+ yr)		<ul style="list-style-type: none"> • Real time remote operation of the experiment 	<ul style="list-style-type: none"> • QoS for latency and reliability 	